XP-002449332

[CHAPTER] 10

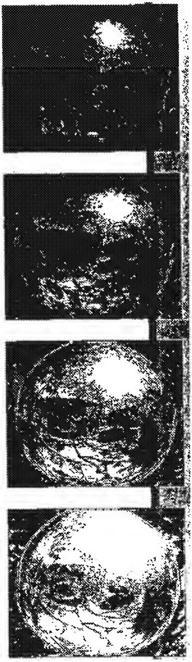

COMPOSITING

FILM AND TELEVISION production companies are continually asked to render scenes that seem to be beyond the limits of their software and hardware. To complete ambitious 3D renderings with the required complexity, quality, and speed, almost all professional productions are rendered in multiple layers or passes, and finished though *compositing*. Compositing is the art of combining multiple images into a unified final scene. Multipass rendering and compositing allow for more efficient rendering, increased creative control, convincing integration with live-action footage, and rapid revisions to your renders.

10.1 RENDERING IN LAYERS

Rendering in layers is the process of rendering different objects in your scene separately, so that a different image is rendered for each layer. The simplest case is to separate animated subjects from background environments. For example, the background layer of the sample scene in Figure 10.1 is the planet surface, and the foreground layer is the spaceship.

To set this up, start with a full scene with all of the objects visible. Assembling all of the objects in one scene will ensure that all layers will match in scale, lighting, and camera angles. After arranging the entire scene, sort your objects into separate layers.

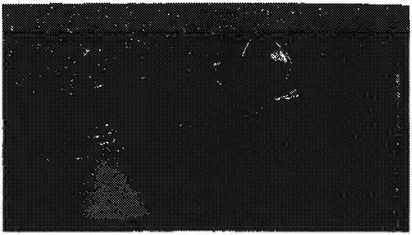

10.4 A sample scene is composited to add the spaceship (from Figure 10.3) over the background environment (from Figure 10.2).

NOTE

Don't have any compositing software yet? See http://3dRender.com/light/compositing/ for an index of all the software companies that supply it.

To render the background layer, shown in Figure 10.2, simply hide the foreground layer objects. If the background layer is not moving, and the camera is not moving, then you need to render only one frame of the background. In your compositing program, the animated foreground can be composited over your static background. If the background is animated, or the camera is moving, then render all of the frames of your background layer.

In a more complex scene, instead of simply hiding the foreground layers during background rendering, you might save a separate version of your scene with the other layers deleted. This can save even more memory and rendering time.

To complete the scene, foreground layers are rendered with the background hidden or deleted. The foreground shown in Figure 10.3 is rendered along with an alpha channel to guide the compositing. The white area of the alpha channel shows where the foreground element will be added. The black areas of the alpha channel indicate areas that will show through to the background layer. For more complex scenes, many foreground layers can be rendered, each with its own alpha.

Why bother to render in layers? Rendering in layers clearly involves more set-up work than rendering all the objects in the scene at once. However, there are several advantages to rendering in layers, some of which can end up saving you valuable time:

- For high-quality character animation, most of your final revisions and rerenders are likely to apply to the character, not the background. In this case, the character can be quickly rerendered as a foreground layer, without rerendering the full scene.

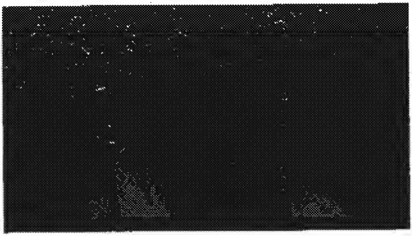

10.2 The background of the sample scene is rendered separately.

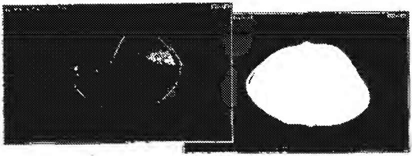

10.3 The foreground layer is another rendering, rendered with an alpha channel to guide compositing.

- For extremely large scenes, rendering in layers makes possible projects that would overload your computer's memory if all objects had to be rendered at once.
- For static background layers, if the camera is not moving, you need only to render one frame to composite behind your entire animated shot.
- For a soft-focused background, you need only to blur the background layer in your compositing program, instead of needing to render a slow depth-of-field effect in 3D.
- For maximum rendering efficiency, different layers can be rendered using different software settings. For example, you might use more antialiasing, motion blur, or raytracing on some important layers. More distant objects, or objects that you plan to blur in post, might be rendered in a manner more optimized for speed.

- Separately rendered elements can be reused in multiple places and times. When compositing layers, you might repeat a flock of birds more than one place in the sky or recycle soft-focused foreground shrubbery between shots.
- To work around bugs, limitations, or incompatibilities in your software, split different effects into different render layers. For example, if you use an effect that does not render through a transparent surface, you could render the effect as a separate layer and then composite the transparent foreground layer over it. If you have created a character's hair via a plug-in that works only with some types of shadows and not others, then you don't have to compromise the lighting of your entire scene if you render the hair in a different layer.

Almost every rendering system has limitations that are best resolved by rendering in layers. In some situations, rendering in layers is more efficient. To achieve even greater efficiency and control, complex professional scenes are usually also rendered in passes.

10.2 Rendering in Passes

Rendering in passes is the process of rendering different attributes of your scene separately. The seven most common types of passes you can render are as follows:

- Beauty pass
- Highlight pass
- Reflection passes
- Shadow passes
- Lighting passes
- Effects passes
- Depth maps

Each of these passes can be created as separate renderings. By combining and adjusting different passes in a compositing program, a scene can be tweaked interactively without being rerendered, and subtle effects can be precisely fine-tuned or matched to a filmed background plate.

10.2.1 Beauty Pass

A *beauty pass* (sometimes called *diffuse pass* or *color pass*) is the main, full-color rendering of your subject, including diffuse illumination, color, and

color maps, as shown in Figure 10.4. A beauty pass usually will *not* include reflections, highlights, or shadows; these are usually rendered as separate passes.

A beauty pass is usually rendered along with an alpha channel of your subject, to be used later for compositing your subject over a background. Different objects in your scene can have different beauty passes if, for example, you want to render different subjects or elements in the environment separately.

To render an object's beauty pass, often a modification to its materials or shaders will be necessary, to block any reflections or specular highlights. Some software may have pass-management features that simplify this process, but editing your subject's materials in order to show only diffuse illumination is something that is possible to achieve in almost any rendering software. A shortcut in many programs is to set your lights to emit only diffuse illumination when rendering your beauty pass (instead of editing each individual material's specularity).

10.2.2 Highlight Pass

Highlight passes (sometimes called *specular passes*) isolate the specular highlights from your objects. You can render highlight passes by turning off any ambient light and setting the object's diffuse shading and color mapping to pure black. The result, as shown in Figure 10.5, will be a rendering of all the specular highlights in the scene, without any other types of shading.

Isolating your specular highlights into a separate pass is optional. Use a separate highlight pass for more important renderings, where you care

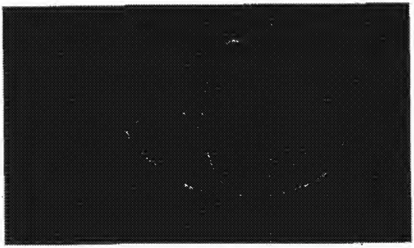

10.4 The beauty pass is rendered without reflections, highlights, or shadows.

10.5 The highlight pass shows only the highlights on the surface, over a black color.

NOTE

LightWave users may want to use Adrian's Nifty Special Buffer Saver from www.1410.com.au/adrian/sbs.htm to help isolate specular highlights, reflections, and many other elements from the buffers created for a rendering.

the most about quality. For less important elements, you might simply include some specularity along with your beauty pass.

Rendering a separate highlight pass allows you more creative control over how the highlights are rendered. In Figure 10.5 a bump map was added to vary and soften the highlights. The bump map was not there in rendering the beauty pass; it was applied only for the highlights. You may also move your lights to different positions if it makes better highlights. Naturally, the lights should come from the same general angle as the lighting is used in the beauty and shadow passes, but there's nothing wrong with cheating a little bit to make a better-looking rendering.

Highlight passes can be composited over your beauty pass with an Add or Screen operation. This way, lighter areas of your highlight pass will brighten your beauty pass, black areas of your highlight pass will have no effect, and no alpha channel is necessary.

During your composite, having highlights as a separate pass will allow control over their color and brightness so that you can adjust the highlights to match the rest of your composited scene. Don't clip large areas of your highlights into pure white. Your highlight pass will work best if highlights run through different shades of gray, which will allow it to look realistic when composited with different intensities.

Separately rendered highlights can also be used to control visual effects, such as glows, added in compositing. Adding a blurred copy of your highlight pass will create glows around your highlights, as shown in Figure 10.6. This way, glows do not take any test renders to adjust, and they can be adjusted in context with the final composite.

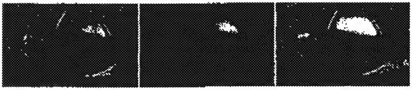

10.6 A composited image (left) can be enhanced with a blurred copy of the highlight pass (middle) to produce a soft glow around the highlights (right).

As discussed in the previous chapter, specular highlights are closely related to reflections, and may be partially or wholly replaced by reflections. If your light sources are visible in a reflection, you might sometimes skip the highlight pass. If your reflection does not include visible reflections of the lights themselves, adding together a highlight pass and a reflection pass can more completely simulate reflected light on the object.

10.2.3 Reflection Passes

A *reflection pass* includes reflections of other objects or the surrounding environment. To isolate reflections, usually all you need to do is turn off ambient, diffuse, and specular shading from a surface so that only reflections appear, as in Figure 10.7.

You do not need any lights to illuminate an object when rendering the reflections on its surface. If your reflections come from reflection maps, you can remove all the lights in your scene when rendering a reflection pass. If raytraced reflections show other objects, some lights may still be needed for the other objects, unless they are constant-shaded or incandescent.

Often, your reflection pass will show reflections of objects from other layers or passes. If you are using reflection maps, this requires that you

10.7 The reflection pass shows only the reflections of the surrounding environment, from a purely reflective object.

have previously rendered a reflection map of your surrounding objects, but you don't need those objects to be present after creating the map. If you are using raytraced reflections, then you need to designate the reflected objects to be visible in reflections (sometimes called *secondary rays*) but not directly visible to primary rays.

NOTE

3D Studio MAX users can use the free Utility Material from Blur Studio (www.blur.com) to render an object's reflections without rendering the original object.

In many cases, you will get the best results by blurring your reflection pass slightly in your compositing program. If you're going to be blurring the reflection pass, you can usually save rendering time by turning off antialiasing. Even if your reflection pass doesn't look perfectly refined when rendered, it will still look good after it's blurred in your compositing program.

10.2.3.1 Compositing Reflections

To composite the passes of a reflective object, a reflection pass is generally dissolved with the underlying beauty pass, before the highlight pass is added on top. A reflection pass may use its own alpha channel, but it can also share the alpha channel from the beauty pass. If you render reflection passes without antialiasing, use the alpha from the beauty pass.

The intensity of the beauty pass, reflection pass, and highlight pass are all reduced when they are blended in the composite, as shown in Figure 10.8. The exact proportions of each are judged visually. As you adjust the passes, compare the highlight pass to the light sources in the scene, and compare the reflection pass to the surrounding environment, to make sure your levels look appropriate to your object and its surroundings. This part of compositing is an art—not a science—and you need to trust your eyes.

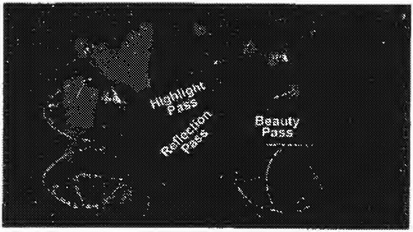

10.8 A balanced composite combines the beauty pass, highlight pass, and reflection pass.

10.2.3.2 Adding Reflections to Real Objects

Sometimes, your 3D object will cast a reflection onto a real-world surface, such as a shiny floor, countertop, or water surface. The cast reflection should be rendered as a separate reflection pass.

If the surface showing the reflection is another 3D object, then that object needs only to be made visible and reflective to render the pass. If the surface showing the reflection is a real-life surface, then you will need to build a 3D model of that surface and position it to align with the live-action plate, as with the grid in Figure 10.9.

For the effect of rippling water, you might add a bump map to a reflective water surface object, which distorts a reflection, as shown in Figure 10.10.

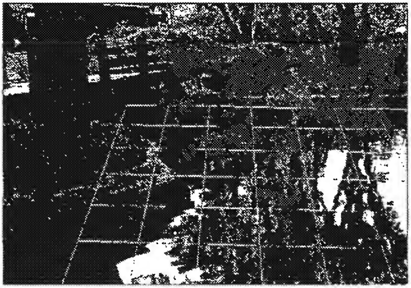

10.9 A grid matches the water surface to receive a reflection.

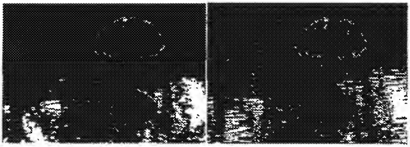

10.10 Bump mapping on a reflection surface distorts the reflection to simulate water ripples.

The alpha channel output by most renderers will default to showing the shape of an entire reflective object, rather than the shape of the reflection itself. For example, the alpha channel could show the shape of the rectangular plane, instead of the reflection that was cast onto the plane. If the alpha channel cannot be fixed in rendering your reflection pass, you may be able to produce a mask by using a luminance key function in your compositing software to isolate the reflection from a black surrounding area. As a last resort, you could render a second reflection pass—reflecting solid white objects against black backgrounds—to use in place of the alpha channel.

When keying the reflection over a surface, adjust it in your compositing program to match the tones of existing reflections on the surface, as in Figure 10.11. Reflections that have too much color and contrast can look unrealistic, so color correction is usually necessary to match any real background sequence.

10.2.4 Shadow Passes

A *shadow pass* is a rendering that shows the locations of shadows in a scene. A shadow pass often appears as a white shadow region against a black background, as in Figure 10.12. Common variations on this include renderings with black shadows against white backgrounds and renderings with the shadow shape embedded in the alpha channel. These variations can all work equally well in compositing.

10.11 A 3D reflection is keyed over water.

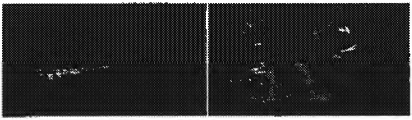

10.12 A shadow pass (left) indicates a ship's shadow cast on the ground and is used to darken the corresponding area of the final composite (right).

To render shadows cast from one object onto another, the object casting the shadows needs to be invisible to primary rays in your rendering but still cast shadows into the scene (this is sometimes called a *shadow object*.) The surface receiving the shadow, such as the ground, still needs to be visible.

To ensure an accurate shadow pass, make sure there is no ambient or incandescent light in the scene. The only light source in the scene should be the one light that is casting the shadow. To further isolate shadows, sometimes a light is used with a negative brightness, as well as a negative shadow brightness, to create white shadows over a uniform black background. This avoids any shading variation in the background, because the background area remains pure black, even beyond the edge of the light's beam.

Shadows often look best when the shadow pass is blurred slightly during compositing. If you are rendering a shadow pass that will be blurred later anyway, you might turn off antialiasing on the shadow pass, turn down shadow samples on an area light, or use a lower resolution shadow map.

NOTE

3D Studio MAX users can use the free Cast Shadows Only Material from Blur Studio (www.blur.com) to render cast shadow passes, and may also use the free Render Layers script from www.gfxcentral.com/bobo/ to help with multiple passes.

10.2.4.1 Attached Shadows

Attached shadows are shadows a surface casts onto itself. For example, a character's nose can cast an attached shadow onto other parts of his face. A bumpy or hilly terrain can also have attached shadows, as shown in Figure 10.13, especially if your light comes from the side.

If you want to control the intensity and color of the shadows separately during your composite, attached shadow passes need to be rendered separately from cast shadows. If you don't want to render an attached shadow pass, attached shadows can sometimes be rendered as a part of the object's beauty pass. In other cases, attached shadows can be skipped entirely. For example, no attached shadows were rendered on the spaceship in the sample scene.

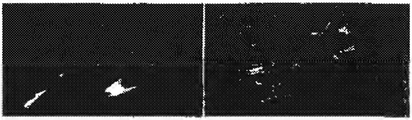

10.13 The attached shadow pass for a terrain is rendered separately (left) and used to darken areas of the final composite (right).

10.2.4.2 Doubled Shadows

A common problem in compositing visual effects shots is doubled shadows. As shown in Figure 10.14, *doubled shadows* occur where a shadow pass further darkens an area that was already shadowed from the same light source. When combining two shadow passes that represent the occlusion of the same light, be sure to composite the passes first in a lighten-only or darken-only mode, and then use the merged shadows to darken the background plate.

Doubled shadows can also be a problem when adding a shadow pass to a shadowed area of a live-action background plate. Especially for exterior scenes where the sun provides a single source for shadows, doubled shadows would look highly unrealistic. The shadow pass must not darken areas that are already darkened by a real shadow. The real shadow needs to be masked out so that your shadow pass darkens only the areas that weren't already in shadow. Your shadow pass should appear to extend the area of the real shadow, extending its area with the same shadow tone.

Adding shadows to a real scene is sometimes a lot of work. Luckily, all that needs to be done in 3D is to render each shadow pass separately, and most of the work is left to the compositing process.

Some 3D artists have the unfortunate habit of trying to render cast shadows into the alpha channel of the same pass as their 3D objects. Because of all the issues that can arise in compositing shadows and all of the adjustments that need to be made to achieve convincing shadow

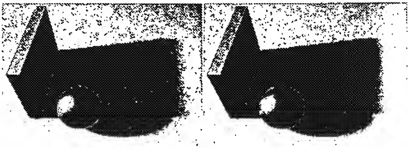

10.14 Poor compositing can result in doubled shadows, with the same light being removed twice from the same area (left, note where the two shadows overlap). Properly composited shadows merge into a unified area where the objects block the same light source (right).

tones, this technique would not actually save time or make compositing any easier.

10.2.5 LIGHTING PASSES

A *lighting pass* is an optional part of multipass rendering that adds a great deal of flexibility and control to the compositing process. Instead of rendering a beauty pass all at once, you could instead render multiple lighting passes, as shown in Figure 10.15. An individual lighting pass shows the influence of one light (or one group of lights) on an element, with other lights hidden or deleted.

The three passes can be added together during compositing. Use the *Add* or *Screen* mode in your compositing program in just the same way that you would composite a highlight pass. During the composite, you can separately tweak the brightness and color of each of the lighting passes to produce different lighting effects, as shown in Figure 10.16.

You probably won't have time to render separate lighting passes from each light on each element in your scene. Separating lighting passes is best reserved for special situations where they will be of the most use during compositing:

- A back light or rim light to an object could be rendered in a separate lighting pass, instead of being a part of your beauty pass. Isolating the rim light could allow the rim light to be separately adjusted when building the final composite.
- Any light that might be blocked by a shadow pass, while allowing other lights to fill in the shadow area, should be rendered as a separate lighting pass.

10.15 Lighting passes were rendered separately for the key light (left), fill light (center), and backlight (right).

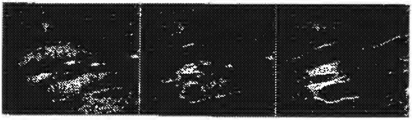

10.16 During compositing, manipulation of the color and intensity of each lighting pass allows lighting adjustments without rerendering.

- Any kind of global illumination, such as radiosity or caustic effects, could be isolated from the direct lighting passes so that the global illumination could be brightened, darkened, or tinted without a recalculation.
- Animated lighting effects are sometimes motivated by a source seen in another pass or live-action plate. For example, a flash of light received from an explosion could be rendered as separate lighting passes. This allows the animated lighting to be precisely matched in color and timing to the element that motivates it.
- Lights with highly saturated colors can be isolated into a separate lighting pass. This allows the hue, intensity, and saturation of the colored light to be edited independently.

When rendering separate lighting passes, try to err on the side of less-saturated light colors. If you need to increase the saturation during the composite, you can, but you wouldn't want to use such a boldly saturated color that the red, green, or blue would be underexposed and difficult to boost.

10.2.6 Effects Passes

Depending on the needs of your project, *effects passes* may sometimes be rendered in a scene. An effects pass is a separate rendering of a visual effect or a mask for a visual effect. An effects pass might be an optical effect, such as a light glow or lens flare, or a particle effect, such as a cloud of smoke or plume of jet exhaust.

10.2.6.1 Optical Effects

Optical effects are phenomena such as lens flares or streaks radiating from a light, which simulate effects that could occur in a camera's lens or film plane. During compositing, optical effects should always be added last, superimposed over other elements.

Optical effects are especially important to isolate as a separate pass. By themselves, they are quick to render. However, you wouldn't want to repeatedly rerender a more complex scene just to see different adjustments to a lens flare!

Optical effects are a popular part of many computer graphics scenes. Visible rays radiating from the star seemed like an obligatory component in making the example scene for this section. However, if you don't want your work to look like a computer graphics demo, you are usually best off to avoid conspicuous optical effects. Rendering optical effects separately at least enables you to diminish or delete the effect later, if you change your mind about its importance in your scene.

10.2.6.2 PARTICLE EFFECTS

The appearance of particles can be modified and enhanced by rendering them separately in their own pass. You can use rendered particles as masks to control different image processing effects, manipulate their color and opacity, and combine them in different ways with the background image.

Figure 10.17 shows a simple cloud of particles. By using the particle cloud as a mask for a distortion effect, the background is distorted by the particles, as though the particles were causing refraction. Finally, the particles are colored green and keyed over the background, behind the ship that emitted them.

When particles are widely distributed throughout a scene, some of them may appear in front of other objects, and some of them behind other objects. There are three main strategies for dealing with this problem:

- Render multiple layers of particles, separating foreground particles, mid-ground particles, and background particles. The particle layers may be separately emitted, or split via clipping planes.
- Render your particles with solid-black copies of other objects, which will block some of the background particles.
- Render particle systems and integrated objects along with depth maps (see the next section) for use in depth-based compositing.

Rendering multiple layers of particles is usually the simplest of these choices, and it allows the most flexible control over particle appearance and compositing.

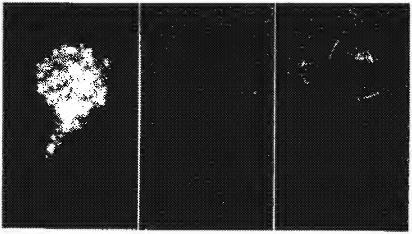

10.17 A simple particle system is rendered as an effects pass over a black background (left) and used as a mask for glass distortion and Gaussian Blur filters applied to a background layer (center) to simulate refractive smoke in the final composite (right).

10.2.7 Depth Maps

A *depth map* (also called *Z-Depth* or a *depth pass*) is a pass that stores depth information at each point in your scene. Some productions use depth maps rendered in a special depth map file format. Other productions use *simulated depth maps,* which are rendered as standard image files just like any other pass but are designed to serve the same purpose as a true depth map.

A depth map is an array of values, measuring the distance from the camera to the closest subject being rendered at each pixel. A depth map of the space scene is shown in Figure 10.18. Brighter shades of gray represent parts of the scene that are closer to the camera.

10.2.7.1 Types of Depth Maps

A true depth map is not an image. The value stored for each pixel location is not a color or shade, but a floating-point accurate number representing the camera-to-subject distance. If a true depth map is converted into a viewable grayscale image, with only 256 levels of gray, then the accuracy is limited, compared to the original distance measurements.

True depth maps output by a renderer do not include antialiasing, because they store only one distance per pixel. You may have to render your images and depth maps at a higher resolution, and then scale down the final composite, in order to achieve correctly antialiased output from a depth-map-based composite. To produce antialiased output from a scene's depth, you can make a rendering that is sometimes called a *simulated depth map*. Simulated depth maps are really the output of a regular rendering. They do not involve floating-point distance measurements. To set up a simulated depth map rendering, all of the objects in your scene

10.18 Brighter tones are closer to the camera in this grayscale representation of a depth map.

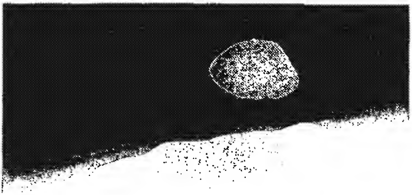

need to be given a flat, constant white color, with no shading. Then, a *depth-fading* or *fog* effect needs to be activated, to fade your scene toward black at greater distances from the camera. The results will be something like Figure 10.18.

10.2.7.2 Uses for Depth Maps

Rendered depth information can be used in different ways during compositing.

If you have two rendered scenes, each of which was rendered with a matching depth map, then a *depth-based composite* can merge the two scenes based on the depth at each pixel. Different objects or particles in the renderings can occlude each other or appear in the foreground or background based on their own distance from the camera.

Depth maps were once used with standalone programs for rendering particles. To composite particle renderings with the objects rendered in another program, depth maps were rendered with both the particles and the objects so that particles could correctly appear in front of or behind different objects. This approach has fallen out of favor because it did not support antialiasing, did not work correctly with transparent objects, and did not allow lighting and shadows to interact between particles and geometry. Particle support built in to the main renderer of most programs has largely replaced this approach.

Another way you can use a depth map is as a mask for any kind of image-processing effect. Figure 10.19 used the rendered depth map (from Figure 10.18 earlier) to mask a blur of the background and to tint the background with a cool, gray color. This enhances the sense of depth

10.19 A depth map is used as a mask for blurring and color-correcting operations that simulate depth.

in the scene through simulating a camera's limited depth of field and also simulates atmospheric perspective as though looking through dust that dulled the color of the distant rocks.

Depth maps are worth rendering for any environment that needs a lot of atmospheric or depth-of-field effects to be simulated during compositing. Simulated depth maps, rendered as a separate pass with depth fading or fog, usually work best for this purpose because their antialiasing will match the other passes.

10.2.8 The Origins of Passes

There is a misconception among some users that rendering in passes is a new technique, only done in certain high-end 3D programs. Rendering in passes has actually been a part of 3D rendering from the first uses of computer graphics in film. In fact, visual effects have been created in passes since before the first 3D effects shots.

The word *pass* originated in *motion-control photography*. Motion-control photography is primarily used in filming miniature models for visual effects. A motion-control camera can repeat the exact same motions several times, moving on a computer-controlled arm or motion base. Each time it repeats its motions through a scene, it's called a *pass*. For example, in filming a model of a spaceship, the first pass through the scene could be a beauty pass, shooting a fully-lit shot of the model. For the next pass, the camera could be loaded with high-contrast film, the lighting could be changed, and the camera could film a matte pass (the equivalent of an alpha channel.) For a third pass, lights might be turned on inside of the model's windows so that a lighting pass could be filmed. Prior to the development of digital compositing systems, all of these passes used to be printed together optically on film.

10.2.9 Pass-Management Features

Any 3D rendering software that is used professionally in film and television productions can be made to render in passes. In most cases, passes are simply created as modified versions of a 3D scene. For example, to isolate a highlight pass, save a version of the scene that has pure black diffuse shading, no ambient light, and renders only specular highlights over a black background. 3D software does not need any special features to support this kind of pass rendering, because all passes are set up manually by the user.

Recently, new pass-management features that can speed and simplify rendering in passes have appeared in some high-end programs, as shown in Figure 10.20. The advantages of having pass-management features

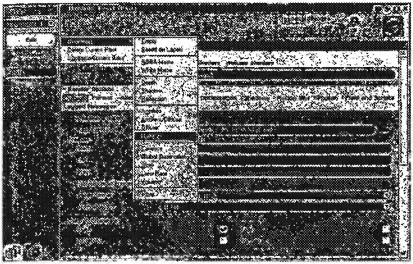

10.20 Pass-management features in SoftimageIXSI streamline the definition of different passes and allow a scene to be instantly switched between pass configurations.

built in to the software are that many pass descriptions can be preconfigured for rapid set-up and that different versions of the scene do not need to be saved separately for each pass. Even after making changes to other aspects of a scene, such as to the animation, all of the passes are still renderable from the same scene.

Another approach to pass management is to render multiple passes and scene layers into one multilayer file, such as RPF (Rich Pixel file Format), the expanded version of the earlier RLA format files output by 3D Studio MAX. If you are sure that the compositing software can make use of the RPF format, then it can be used to convey images, Z-Depth, and other information that allows for effects and lighting changes, all in a single file.

Pass-management features in 3D software are a case of the developers following their users. Advanced 3D effects have always been accomplished by layering different passes over a background plate, even before software developers started to notice and build in compositing-related rendering features.

10.3 Lighting to Match Background Plates

A *background plate* is usually a sequence of frames digitized from live-action film or video, into which you will add your computer graphics elements. In some projects, your background plate could also be a matte

painting, a previously rendered 3D scene, or a composite of multiple elements.

Many 3D programs have an option to view background plates within an animation window (sometimes called an *image plane* or *rotoscope background*), showing how your subject will be aligned with the background. After your 3D scene is aligned with the camera angles from the real-world shot, you face the challenge of matching the lighting from the real-world environment. Matching the direction, color, and tone of the light sources in the real scene is essential to integrating your 3D rendered passes with the photographed background plate.

10.3.1 Reference Balls and Light Probes

A set of reflective and matte balls can be ideal reference objects to help measure the position and color of lights on a location. Mirrored balls are sold as lawn ornaments and as housings for ceiling-mounted security cameras. For a ball with a matte finish, plaster is ideal, but a Styrofoam ball from a craft store could be more portable and affordable. You may need to paint the ball gray and attach a piece of wire to hold it in place.

10.3.1.1 Matte Balls

A picture showing the matte ball in a lighting environment, as shown in Figure 10.21, can be great for choosing the color of light reaching your subject from each direction. Ideally, this image should be shot with the same camera as your final background plate and digitized in the same session with the same color correction.

For the most accurate color matches, bring the ball image into a paint program and pick specific RGB color values from the image of the ball, as shown in Figure 10.22. These RGB colors can then be assigned directly as colors for your lights from corresponding directions.

When developing the scene's lighting, you can import the ball image as a background image in your 3D program and create a 3D sphere in front of the image. Using your 3D sphere as a reference, adjust infinite or directional lights from each direction to make the shading of the 3D sphere match the shading of the ball in the background plate.

Studying the colors reaching a point in a real-world environment is a great exercise for anyone working in 3D lighting. Even if you don't need to match the lighting of a background plate right now, this process could be worth trying a few times, just to get a better feel for the colors of real-world lights.

10.21 A matte-finished gray ball is positioned between a fire and window to probe different colors in the scene.

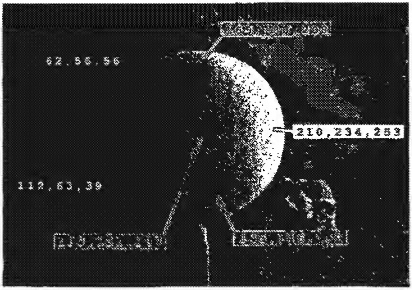

10.22 RGB colors chosen from the reference ball will give an accurate color match to 3D lights.

10.3.1.2 Mirror Balls

A picture of a reflective ball in an environment helps you more precisely determine the angle and relative brightness of each light, and guides you in creating highlights and reflections for your object (see Figure 10.23). As you know, reflections are view-dependent, so it is vital that the reflective ball be shot from the same camera position as your final background plate.

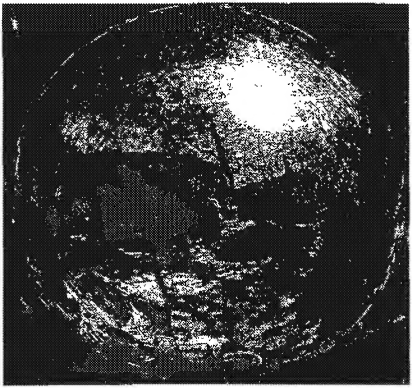

10.23 A mirror ball captures a reflected image of the surrounding environment for reference or for use as a reflection map.

As with the matte ball image, the reflective ball image can be brought into your 3D program. If you make a shiny 3D sphere, you should be able to see highlights from your brighter lights and match these to the highlights in the reflective ball.

An added bonus to having a picture of a reflective sphere in your environment is that you can use it to develop a reflection map for your object, as shown in Figure 10.24. In many programs, the best way to project your acquired image is as a planar projection onto the side of a large sphere surrounding your 3D scene; make the sphere render reflections only.

10.3.1.3 Light Probe Images

The traditional approach to matching natural light from a real environment is to use an array of infinite or directional lights, as described earlier. After each light is properly adjusted to match the color and brightness from each direction, this approach can produce realistic renderings and seamless lighting matches.

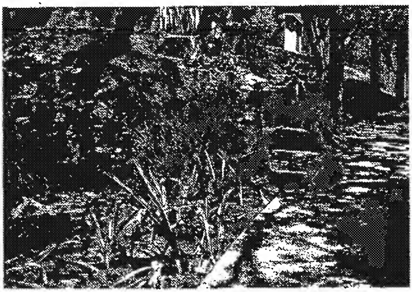

10.24 The captured reflection map matches a 3D object's reflections to the environment.

Another approach to re-creating a real-world lighting environment is to use a *light probe image* recorded on the same location as your background plate. A light probe image is an image that captures the lighting from all angles around a subject, which can be created by photographing a reflective ball at the shooting location. The light probe image can then be used to illuminate objects with all the recorded colors and tones of real-world light. From only one light probe image, illumination will reach 3D objects from all angles, as though the light source were a giant sphere wrapped around the entire 3D scene.

Unlike ordinary photographs of a reflective ball, light probe images are *high dynamic range (HDR) images*, meaning that they can capture an exposure latitude greatly exceeding the range of one visible image. To photograph light probe images, cameras are programmed to shoot a series of images at different exposure settings, exposing for the brightest light sources all the way down to the darkest, as shown in Figure 10.25. Without using HDR images, all of the brighter lights in a scene might appear clipped as pure white highlights, with no record of their relative brightness or color. By using HDR images, a light probe image can accurately record the color and relative brightness of every light source.

NOTE:

Information about High Dynamic Range images can be found on the Internet at www.cs.berkeley.edu/~debevec/ and flatlux.berkeley.edu/mkhdr/. The first commercial software to support illumination with HDR images is NewTek's LightWave 3D (www.newtek.com).

10.3.2 Other Approaches

You can't always use probes and reflective balls on the set of every production. You can't expect every production to stop and wait for you to

10.25 A High Dynamic Range (HDR) image will indicate accurate colors at multiple levels of exposure.

set up special reference shots. Sometimes you won't even be able to visit the location where background plates are photographed.

Even if you do get to measure the lighting with different kinds of balls, the lighting in the scene may change without being remeasured. Balls in one location in a scene also may fail to give you the information you need about lighting in another point—you'd need an infinite number of probes to fully measure the light at every point in space.

If you can go to the set or shooting location, you can use other techniques to assist in matching the lighting:

- **Bring a camera to the set.** Take reference pictures of the set and the lighting around it. Take flat-on pictures of walls or floors for possible use in texture mapping. Take wide-angle or panoramic pictures to create reflection maps.
- **Bring a measuring tape to the set.** Ask early in the production if you can have blueprints to the set, but don't trust the original plans to be accurate. Bring a measuring tape and record enough information so that you could build an accurate 3D model of the set if necessary.
- **Watch for changes during the course of the production.** In a studio, lights are adjusted, and even the walls and set pieces are moved between shots. Outside, the weather and time of day create more changes.

If you cannot go to the shooting location, or your background plate comes from stock footage or other sources, you still can match the lighting using other techniques:

- **Study shadows in the background plate.** When you have matched their length and direction in 3D, your lights will be in the right places.
- **Use an object in the background plate to find light colors.** Try to find a white or gray object in the background plate from which you can pick RGB values.
- **Try to find reference objects in the background plate that can be made into a matching 3D model.** By aligning the 3D model with the real object, you can compare how illumination and highlights hit your 3D model until it receives the same illumination as the background plate.

Every production will create different challenges, but with this basic set of tricks, you should be able to match the lights from any background plate.

10.4 Working with a Compositor

As a 3D artist in a professional environment, you might not do all of your own compositing. Often there will be a full-time compositor assembling effects shots that contain your 3D rendered elements, as well as live-action images and other types of plates.

In commercial production facilities, it makes good economic sense to finish your 3D work and move a project into compositing as soon as possible, making changes and experiments in the compositing stage. 3D-rendering projects are often billed on a per-project basis, while compositing services are billed per hour of use.

The speed of the compositing process is generally fast enough that clients can supervise compositing sessions, and request and evaluate changes as the work is done. This speed and degree of client involvement is not yet possible in 3D (thankfully!), making clients happier to work and request changes in the 2D compositing stage, even if they are paying by the hour.

Give more flexibility to the compositor, and a shot is less likely to need rerendering in 3D. Here are a few things you can do in rendering to give the compositor more flexibility include the following:

- **Keep animated objects within the frame.** There could be times when you need to reposition your object within the frame, to re-animate its position or track it to a 2D camera move. For any element that a compositor might want to reposition, try to frame your 3D shot so that your animated subject is not cropped off at the edge of the frame. Even if the compositor needs to scale the object up, it is better than losing a part of the model that should have been visible.
- **Render "clean" footage, and save post effects for compositing.** Most 3D programs have options to add effects such as glows, blurs, film grain, and lens flares to your 3D rendering. In most cases it is best to render without these effects and instead add them in a compositing program. It is not cost-effective to redo a raytracing just to test different levels of film grain, when the film grain could have been added to the rendered image in 2D more quickly and interactively.
- **Err on the side of motion that is too slow if speed changes might be needed to an element.** Motion can be sped up more easily than slowed down. To speed up footage, frames can simply be thrown away or dissolved together. To slow down footage, new

Note

For more information on compositing in general, check out the books The Art and Science of Digital Compositing by Ron Brinkmann (Morgan Kaufmann) or Digital Compositing In Depth by Doug Kelly (Coriolis). New Riders tells me their own [digital] Compositing & Post will be out in 2001—until then, though...

frames would need to be created, often by repeating existing frames, which can make motion jerky.

- **Use your full exposure latitude.** If you didn't read Chapter 7, "Exposure," before this chapter, please read it now! Vital issues in compositing are matching the black levels, colors, and brightness of the background plate. Compositors will have the most flexibility if they use the full exposure latitude available, so that banding artifacts don't appear during color correction.
- **Avoid clipping.** Areas of your rendering that are overexposed to pure white or underexposed to pure black lack image information and cannot be effectively brightened or darkened during compositing. Keep your shading within a range of tones that will allow for further adjustment.
- **Render more passes.** As a general statement, a scene rendered in more separate passes is easier to manipulate than a scene rendered in only a few passes. For example, if a light source is rendered in a separate lighting pass, then that light can be brightened, dimmed, or tinted by the compositor without rerendering in 3D.

Compositing is a complex craft in its own right and worthy of study by any digital artist. Ideally, you should try to get your first compositing experience working on student, personal, or low-budget projects, where a 3D artist becomes a jack-of-all-trades. Even if you are doing your own compositing on a personal computer, compositing can grow into a great timesaver, a powerful problemsolver, and an important part of your creative process.